

**METHOD AND APPARATUS FOR INCREASING AN AMOUNT OF  
MEMORY ON DEMAND WHEN MONITORING REMOTE MIRRORING  
PERFORMANCE**

5    **BACKGROUND OF THE INVENTION**

The present invention relates generally to techniques for implementing the remote mirroring function. More particularly the present invention relates to a method and apparatus for increasing an amount of memory on demand in a storage system implementing the remote mirroring function when the 10 mirrored pair of volumes of the storage system is operated in the asynchronous mode.

Conventionally, there have been two types of approaches to storage-based volume replication, namely local replication and remote (copy) replication. Both technologies mirror files, file systems, or volumes without 15 using host CPU power. When a host conducts a host input/output (I/O) such as a write I/O of data to a primary volume (PV) of a storage system, the storage system automatically copies the data to a replication (secondary) volume (SV). This mechanism ensures that PV and SV are identical.

Local replication duplicates the primary volume within a first storage 20 system, so that when the host writes data to the PV, the first storage system also stores the data to a local secondary volume (LSV). Local replication is typically used for taking backups.

Remote replication duplicates volumes across two or more storage systems so that when the host writes data to PV, the first storage system 25 transfers the data through paths, such as ESCON, Fibre Channel, T3, and/or IP networks, to at least one second storage system for storage in a remote secondary volume (RSV) included therein. Remote replication is typically

used to enable the recovery of data from disasters, such as earthquake, flood, fire, and the like. Even if the first storage system or the whole data center at the primary site is damaged by a disaster, data is unaffected at the secondary site and business can be resumed quickly.

5        There are at least two modes of transferring data to implement remote mirroring between local and remote storage systems, namely synchronous mode and asynchronous mode. In the synchronous mode, all write I/O's to the PV of the first storage system are mirrored at the RSV of the second storage system. In the asynchronous mode, in response to a write I/O, the first

10      storage system completes the write I/O and then asynchronously transfers the write data to the second storage system for storage on the RSV. Specifically, the write data to be copied to the RSV of the second storage system is temporarily stored in a queuing area, such as cache memory, disk storage, Non-Volatile Random Access Memory (NVRAM) etc. The write data is

15      retrieved from the queuing area and then stored in the RSV of the second storage system.

          The queuing area where the write data is temporarily stored may run out if the write I/O rate (write I/O's per second) to the primary volume of the first (local) storage system is much higher than the throughput for copying the

20      write data from the queuing area to the remote secondary volume of the second (remote) storage system. Once the unused area of the queuing area becomes less than a certain percentage of total area of the queuing area, the storage system mirroring function will initiate a throttling function to curb the new write I/O's. Thus, the performance of the storage system with respect to

25      write I/O's degrade over time.

A current proposed solution is that if a customer experiences a degradation in performance with respect to write I/O's, then the customer calls a customer support center which analyzes the cause of the problem and may increase the capacity of the queuing area if necessary. The total turn around 5 time for this proposed solution if long can negatively affect customer satisfaction. This negative affect on customer satisfaction is more likely to occur since the turn around time is influenced by many human interactions that are necessary.

## 10 SUMMARY OF THE INVENTION

The present invention provides a method and apparatus for increasing an amount of memory on demand in a storage system implementing the remote mirroring function when the mirrored pair of volumes of the storage system is operated in the asynchronous mode.

15 According to the present invention a method and storage system is provided for increasing an amount of memory on demand in the storage system implementing a remote mirroring function. The storage system includes first and second storage subsystems which are connected to each other via a path and the first storage subsystem is connected to a host. A 20 primary volume is provided to the first storage subsystem and a remote secondary volume is provided to the second storage subsystem. The remote secondary volume is a copied volume of the primary volume and the primary volume and the remote secondary volume are operated in the asynchronous mode. A memory is provided in the second storage subsystem for temporarily 25 storing data transferred to the second storage subsystem from the first

storage subsystem in response to a write input/output (I/O) issued by the host to write data in the primary volume. Data temporarily stored in the memory is retrieved and stored in the remote secondary volume. An unused area of the memory is monitored and the amount of the memory is increased if the 5 unused area becomes less than a predetermined amount.

## BRIEF DESCRIPTION OF THE DRAWINGS

The foregoing and a better understanding of the present invention will become apparent from the following detailed description of example 10 embodiments and the claims when read in connection with the accompanying drawings, all forming a part of the disclosure of this invention. While the foregoing and following written and illustrated disclosure focuses on disclosing example embodiments of the invention, it should be clearly understood that the same is by way of illustration and example only and the invention is not limited thereto, wherein in the following brief description of the drawings:

Fig. 1 illustrates the configuration of a system for increasing an amount of memory on demand in a storage system according to the present invention;

Fig. 2 illustrates an example of a disk management table for managing disk drives to implement the present invention;

20 Fig. 3 illustrates an example of a cache management table for managing cache memory to implement the present invention;

Fig. 4 illustrates an example of a monitoring result table for storing the results of monitoring the resources of a storage system according to the present invention;

Fig. 5 is a flowchart illustrating an example of a process for increasing an amount of memory on demand in a storage system according to the present invention; and

Fig. 6 is a flowchart illustrating an example of the details of processing 5 performed to implement Step 503 of Fig. 5 according to the present invention.

#### DETAILED DESCRIPTION OF THE INVENTION

As per Fig. 1, the present invention provides a storage system which includes a first storage subsystem 101 and a second storage subsystem 102 10 that are connected to each other by a path 105 and implements a remote mirroring function. The method and storage system of the present invention as provided increases an amount of memory on demand in the storage system as needed depending on the state of a queuing area.

As per Fig. 1, the first storage subsystem 101 includes a primary 15 volume (P) 103, a storage controller 106-1 which includes a cache management table 190-1, a disk management table 191-1 and a resource/license manager 195-1, and a queuing area 196-1. The first storage subsystem 101 is connected to a host 140 and to a network 180.

Further, as per Fig. 1, the second storage subsystem 102 includes a 20 remote secondary volume (S) 104, a storage controller 106-2 which includes a cache management table 190-2, a disk management table 191-2 and a resource/license manager 195-2, and a queuing area 196-2. The remote secondary volume (S) 104 is a copied volume of the primary volume (P) 103 and the primary volume (P) 103 and the remote secondary volume (S) 104 25 are operated in the asynchronous mode. The second storage subsystem 102

is connected to hosts 107 and 108. The second storage subsystem 102 and the hosts 107 and 108 are connected to the network 180 via a local area network (LAN) 197.

As illustrated in Fig. 1, the queuing area 196-1 and 196-2 of each storage subsystem 101 and 102 can for example be cache memory 1960-1, 1960-2, disk drives 1961-1, 1961-2, or memory devices 1962-1, 1962-2. Any of these devices would be acceptable and of course such devices could, for example, be pre-installed and configured so as to permit expansion of its storage capacity.

According to Fig. 1 a management center system 150 is provided for managing the storage resources of the storage system and various licensing matters with respect to the storage system including various licensing keys as will be described below. The management center system 150 is connected to the network 180 and thereby to the first storage subsystem 101 and the host 140 via the network 180 and to the second storage subsystem 102 and the hosts 107 and 108 via the network 180 and the LAN 197. The management center system 150 includes a storage 198 which serves as a repository of information collected from the resource/license managers 195-1 and 195-2, each of which manages the resources of its respective storage subsystem 101, 102 based on the cache management table 190 and the disk management table 191.

Fig. 2 illustrates an example of a disk management table 191 for managing disk drives to implement the present invention. The disk management table 191 includes a plurality of entries, wherein each entry 25 includes at least a device number field 201, an installation status field 202, a

license key status field 203 and a size field 204. Thus, as per Fig. 2, for example, a first entry of the disk management table 191 indicates in the device number field 201 a disk drive 1, in the installation status field 202 that disk drive 1 is installed, in the license key status field 203 that the license key

5 for disk drive 1 is installed, and in the size field 204 that the size of disk drive 1 is 500 Giga bytes. Another entry of disk management table 191, for example the third entry, is similar to the first entry with the exception that the third entry concerns disk drive 3, that the license key for disk drive 3 is not installed and that the size of the disk drive 3 is 250 Giga bytes. Yet another 10 entry of disk management table 191, for example the fifth entry, is different from the other entries in that the fifth entry concerns disk drive 5 and that disk drive 5 is not installed. According to the present invention a license key must be installed in order to use a corresponding disk drive and of course a disk drive can not be used if it has not been installed.

15 Thus, according to the present invention disk management table 191 indicates whether or not a disk drive has been installed, whether or not a license key for the disk drive has been installed and the size of the disk drive if installed. Therefore, the disk management table 191 tracks information regarding the status of respective disk drives, thereby permitting decisions to

20 be made so as to properly manage the resources of the storage system particularly with regard to the queuing area 196.

Fig. 3 illustrates an example of a cache management table 190 for managing cache memory to implement the present invention. As per Fig. 3, the cache management table 190 includes a plurality of entries, wherein each 25 entry includes at least a slot number field 301, an installation status field 302,

a license key status field 303 and a size field 304. Thus, as per Fig. 3, for example, a first entry of the cache management table 190 indicates in the slot number field 301 a memory device 1, in the installation status field 302 that memory device 1 is installed, in the license key status field 303 that the

5 license key for memory device 1 is installed, and in the size field 304 that the size of memory device 1 is 2 Giga bytes. Another entry of cache management table 190, for example the second entry, is similar to the first entry with the exception that the second entry concerns memory device 2, that the license key for memory device 2 is not installed and that the size of the

10 memory device 2 is 2 Giga bytes. Yet another entry of cache management table 190, for example the third entry, is different from the other entries in that the third entry concerns memory device 3 and that memory device 3 is not installed. According to the present invention a license key must be installed in order to use a corresponding memory device and of course a memory device

15 can not be used if it has not been installed.

Thus, according to the present invention cache management table 190 indicates whether or not a memory device has been installed, whether or not a license key for the memory device has been installed and the size of the memory device if installed. Therefore, the cache management table 190

20 tracks information regarding the status of respective memory devices, thereby permitting decisions to be made so as to properly manage the resources of the storage system particularly with regard to the queuing area 196.

Fig. 4 illustrates an example of a monitoring results table 400 for storing the results of monitoring the resources of a storage system particularly

25 with regard to the queuing area 196 according to the present invention. The

monitoring results table 400, for example, is stored in either the storage 198 of the management center system 150 or in the resource/license manager 195 included in each of the storage subsystems 101 and 102. The monitoring results table 400 includes a plurality of entries wherein each entry includes a

5 device identification (ID) field 401 for identifying a device number of a device being used in the queuing area 196, a device field 402 which identifies the type of device either a cache or disk drive being used in the queuing area 196, a configuration field 403 for indicating an amount of the total capacity of the queuing area the device represents according to the particular 10 configuration, a capacity field 404 which indicates the capacity of the device and a current usage field 405 for indicating the amount of the capacity of the device being used.

Thus, for example, as illustrated in Fig. 4, the first entry of the monitoring results table 400 in the device ID field 401 identifies device 1, in 15 the device field 402 identifies device 1 as a cache, in configuration field 403 indicates that the cache of device 1 represent 70% of the total capacity available, in the capacity field 404 indicates device 1 has a capacity of 32 Giga bytes, and in the current usage field indicates the current usage of the cache amounts to 30% of the available capacity.

20 As per the present invention, the management center system 150 or the resource/license manager 195 makes use of the information stored in the monitoring results table 400 so as to determine whether an action should be taken or not. Also, the monitoring results table 400 can be used to permit historical analysis of the usage made of the devices forming the queuing area

196 so as to determine whether a particular trend is exhibited to which a counter process or procedure can be implemented.

Fig. 5 illustrates a flowchart of an example of a process executed by either the management center system 150 or the storage subsystems 101 or

5 102. It should be further noted that each of the steps illustrated in the flowchart of Fig. 5 can, for example, be implemented by one or more instructions of a computer program stored in a storage medium such as the hard drive of a computer system, a floppy disk, CD ROM, electrically erasable programmable read only memory (EEPROM), etc.

10 As per Fig. 5, the process starts with the user defining, setting the event or condition to be detected and the specific action to be performed when the event or condition has been detected (step 501). The event or condition to be detected could, for example, include detecting when the current usage of the queuing area 196 exceeds a predetermined limit. This 15 limit could be set at any level (60%, 80%, etc.) so long as it is appropriate to ensure efficient operation of the storage system. Information regarding the set event or condition is periodically checked (monitored) so as to detect whether the event or condition has occurred (step 502). The information to be checked could, for example, include the entries in the monitoring results table 20 400. If the event or condition has not occurred then the system continues monitoring the information (step 505). If the event or condition has occurred then the action set by the user corresponding to the event or condition is performed (step 503). The details of step 503 are illustrated in Fig. 6 and will be discussed below. After the preset action has been performed a result of 25 the action is output, for example, in the form of a notice or charge (step 504).

As per the above, Fig. 6 illustrates a flowchart of an example of the details of processing performed to implement step 503 of Fig. 5. As per Fig. 6 in the action of step 503 the preset user defined action including, for example, the management center alerting the user that the preset event or condition

- 5 has occurred by email, pager, etc., is performed (step 5031). Thereafter, it is determined whether the additional capacity (memory) needed by the queuing area 196 is to be automatically provisioned or not (step 5032). If the additional memory is not to be automatically provisioned then the system waits for a confirmation as to whether the user defined action is to be
- 10 conducted or not (steps 5033 and 5034). This confirmation could, for example, be provided by the management center system 150 or the user of the storage system being monitored. If the confirmation has not been received then step 5033 is conducted again. If the confirmation has been received then a new license key is sent to the storage subsystem requiring an
- 15 increase in capacity of its queuing area 196 or the license key is removed from the storage subsystem (step 5035). Information indicating a result of the action taken, for example, installing a disk drive to the queuing area is sent to either the resource/license manager 195 or the management center system 150 (step 5036). Thereafter, the information indicating the result sent
- 20 to either the resource/license manager 195 of the management center system 150 is stored in the monitoring result table 400 (step 5037).

It should be noted that step 503 of Fig. 5, the details of which are illustrated in Fig. 6 could, for example, be performed by the management center system 150 while steps 501, 502, 504 and 505 could, for example, be

- 25 performed in either the management center system 150 or a storage

subsystem 101 or 102. Further, it should be noted that the storage 198 could, for example, also serve as a repository of collected information including, for example, customer definition regarding features of the customer's storage subsystems and location information regarding the location of the customer's

5 storage subsystems. This information is consulted as needed by the management center system 150 when automatic provisioning is not performed. In such a situation, a service person receives such information from the management center system by email, pager or the like and uses such information to service the customer's storage subsystem at the

10 customer's location upon detecting or being given notice that an unused area of the queuing area is less than a predetermined level.

Therefore, based on the above, the present invention provides a method and apparatus for increasing an amount of memory on demand in a storage system implementing a remote mirroring function. According to the

15 present invention, the storage system includes first and second storage subsystems which are connected to each other via a path wherein the first storage subsystem is connected to a host. As per the present invention, a primary volume is provided to the first storage subsystem and remote secondary volume is provided to the second storage subsystem, the remote

20 secondary volume being a copied volume of the primary volume which is operated in a synchronous mode with respect to the remote secondary volume. As per the present invention, a memory such as cache memory, disk drives, memory devices or the like are provided in the second storage subsystem for temporarily storing data transferred to the second storage

25 subsystem from the first storage subsystem in response to a write I/O issued

by the host to write data in the primary volume. Data temporarily stored in the memory is retrieved and stored in the remote secondary volume. However, unique according to the present invention is that an unused area of the memory is monitored and the amount of the memory is increased if the 5 unused area becomes less than a predetermined amount. Thus, the present invention provides for quick and efficient provisioning of additional memory in the queuing area of a storage system on demand.

While the invention has been described in terms of its preferred embodiments, it should be understood that numerous modifications may be 10 made thereto without departing from the spirit and scope of the present invention. It is intended that all such modifications fall within the scope of the appended claims.